Scientific Computing

OLCF
Oak Ridge Leadership Computing Facility

Ricky Kendall

Group Leader, Scientific Computing Group

…as told by

Bronson Messer

Scientific Computing

Scientific Computing facilitates the delivery of leadership science by partnering with users to effectively use computational science, visualization, and workflow technologies on OLCF resources to:

Port, tune, augment, and develop current and future applications at scale	Provide visualizations to present scientific results and augment discovery processes	Automate the scientific computational method

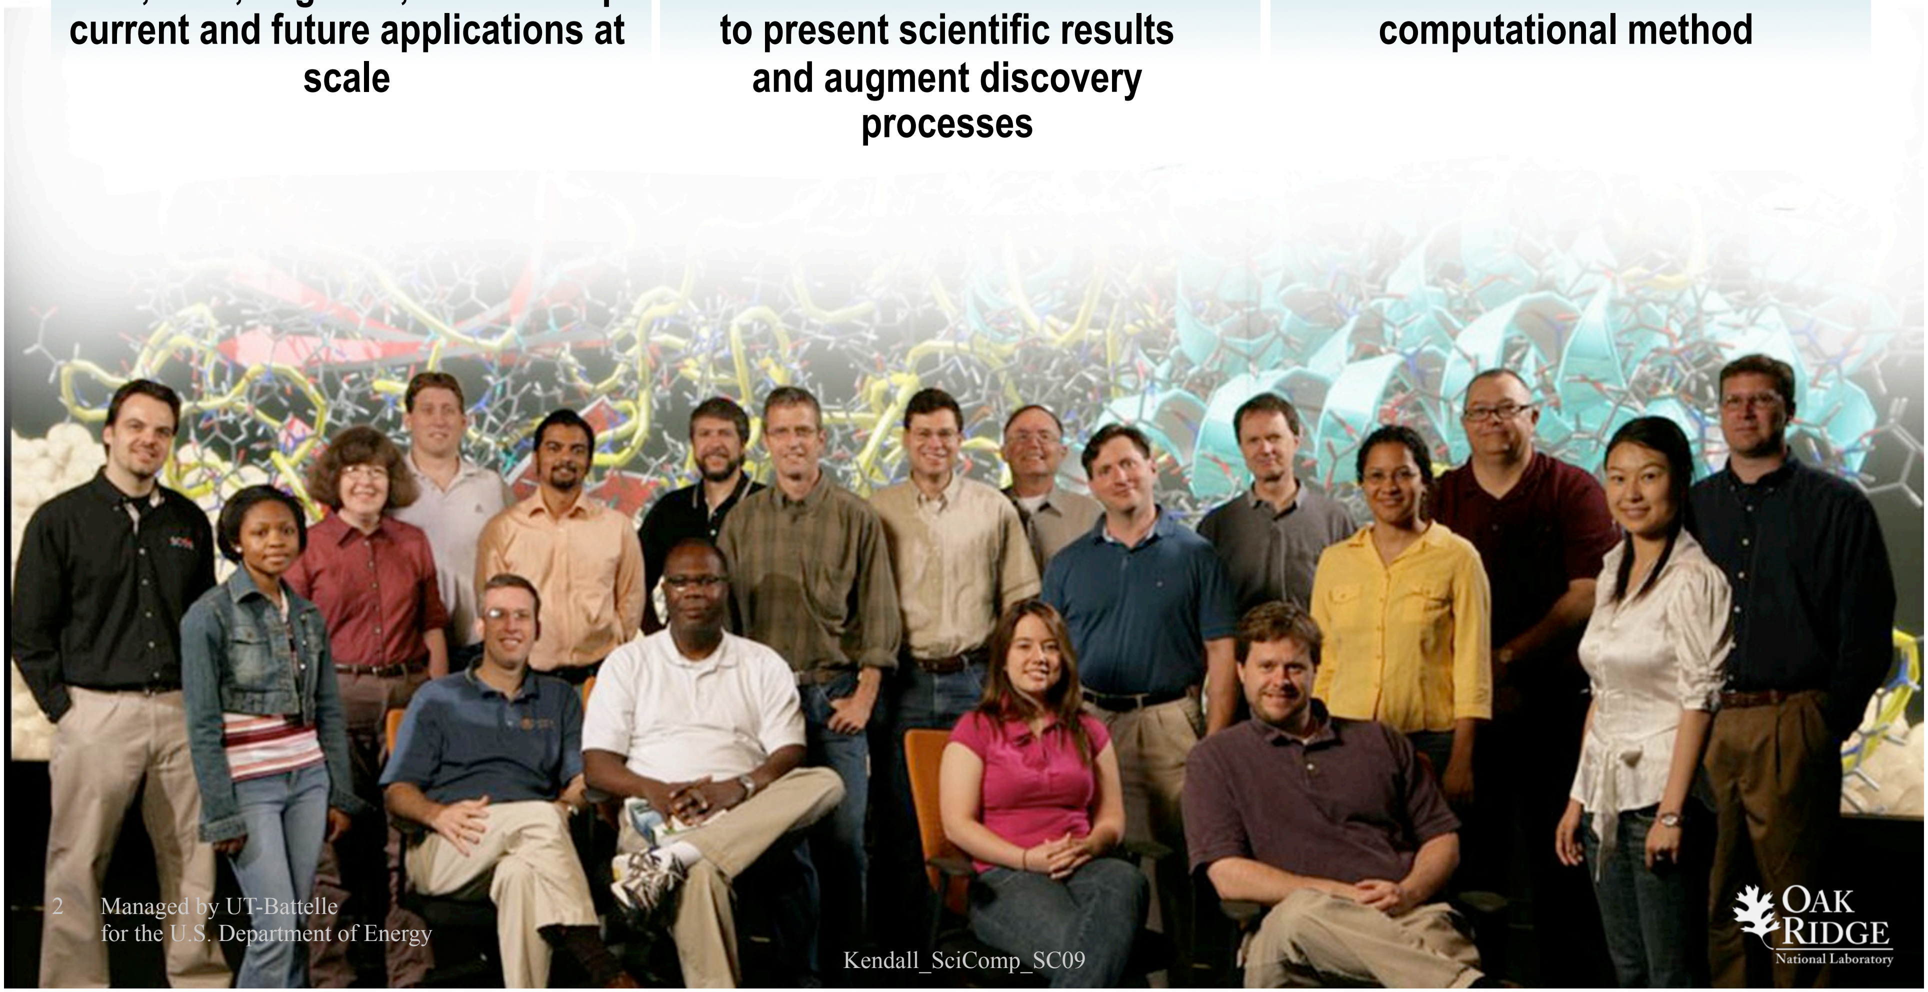

Operations Council
W. McCrosky, Finance Officer
H. George, HR Rep.
K. Carter, Recruiting
S. Milliken, Facility Mgmt.
D. Edds, ES&H Officer
R. Adamson, M. Disney, Cyber Security

National Center for Computational Sciences
Oak Ridge Leadership Computing Facility
J. Hack, Director
A. Bland, OLCF Project Director
L. Gregg, Division Secretary

Advisory Committee
J. Dongarra
T. Dunning
S. Karin
D. Reed
J. Tomkins

Director of Science
B. Messer (Acting)

Industrial Partnerships
S. Tichenor

Director of Operations
J. Rogers

INCITE Program
J. White

Chief Technology Officer
A. Geist

OLCF System Architect
S. Poole

Deputy Project Director
K. Boudwin

B. Hammontree, Site Preparation
G. Shipman, File Sys. Acquisition, Dev. & Commissioning
J. Rogers, Computer Acquisition
R. Kendall, Pre-commissioning & Acceptance Test Dev.
A. Baker, Commissioning
D. Hudson, Project Management
A. Barker, Training & Support Development

Cray Supercomputing Center of Excellence
J. Levesque, N. Wichmann, J. Larkin
D. Kiefer, L. DeRose, K. Seymour

Application Performance Tools[5]
R. Graham
T. Darland

M. Baker
J. Dobson
O. Hernandez
S. Hodson
C. Hsu
J. Hursey[7]
T. Ilsche
T. Jones
C. Kartsaklis
G. Koenig
J. Kuehn
J. Ladd
T. Mintz
B. Settlemyer
P. Shamis
M. Gorentla Venkata

Scientific Computing
R. Kendall
A. Fields

M. Abbasi[7]
S. Ahern[#]
V. Anantharaj
E. Apra[5]
D. Banks[3]
G. Bisht
M. Brown
J. Daniel
M. Eisenbach
M. Fahey
J. Gergel[5]
R. Hartman-Baker
J. Hursey[7]
W. Joubert[#]
S. Klasky[#]
R. Kumar[7]
J. Logan[7]
A. Lopez-Bezanilla[7]
C. Ma[1]
M. Matheson
R. Mills[5]
B. Mintz[7]
H. Nam
M. Norman
G.Ostrouchov[5]
N. Podhorszki
D. Pugmire
R. Sisneros[7]
R. Sankaran
S. Su[7]
R. Tchoua
A. Tharrington[#]
R. Toedte

User Assistance and Outreach
A. Barker
S. Mowery

J. Buchanan
A. Carlyle
C. Fuson
E. Gedenk[1]
B. Gajus[5]
M. Griffith
S. Hempfling
J. Hines[#]
S. Jones[5]
D. Levy[5]
M. Miller
L. Rael
B. Renaud
C. Rockett[1]
D. Rose
A. Simpson
J. Smith
B. Whitten
L. Williams[5]

Technology Integration
G. Shipman
S. Mowery

S. Atchley
T. Barron
D. Dillow
D. Fuller
R. Gunasekaran
J. Harney[7]
S. Hicks[5]
Y. Kim
K. Matney
R. Miller
S. Oral
B. Settlemyer[5]
D. Steinert
J. Simmons
V. Tipparaju[5]
S. Vazhkudai[5]
F. Wang
V. White

High-Performance Computing Operations
A. Baker
S. Allen

R. Adamson
J. Anderson
M. Bast
J. Becklehimer[4]
J. Breazeale[6]
J. Brown[6]
M. Disney
A. Enger[4]
C. England
J. Evanko[4]
A. Funk[4]
D. Garman[4]
D. Giles
M. Hermanson
J. Hill
S. Koch
H. Kuehn
C. Layton
C. Leach[6]
J. Lewis[4]
D. Leverman
D. Londo[4]
J. Lothian
D. Maxwell[@]
M. McNamara[4]
J. Miller[6]
D. Pelfrey
G. Phipps, Jr.[6]
R. Ray
S. Shpanskiy
C. St. Pierre
B. Tennessen[4]
K. Thach
J. Trucks
J. Walsh[4]
T. Watts[4]
S. White
C. Willis[4]
T. Wilson[6]

[1]Student
[2]Post Graduate
[3]JICS
[4]Cray, Inc.
[5]Matrixed
[6]Subcontract
[7] Post Doc
*Acting
[#] Task Lead
[@] Technical Coordinator

Visualization and data analytics

Visualization

Once users have completed their runs, the Visualization task group helps them make sense of the sometimes overwhelming amount of information they generate

- Viewing at a 30 ft × 8 ft PowerWall
- Cluster with GPUs for remote visualization

End-to-End Solutions

Researchers must analyze, organize, and transfer an enormous quantity of data. The End-to-End task group streamlines the workflow for system users so that their time is not eaten up by slow and repetitive chores

- Automate routine activities (e.g., job monitoring at multiple sites)

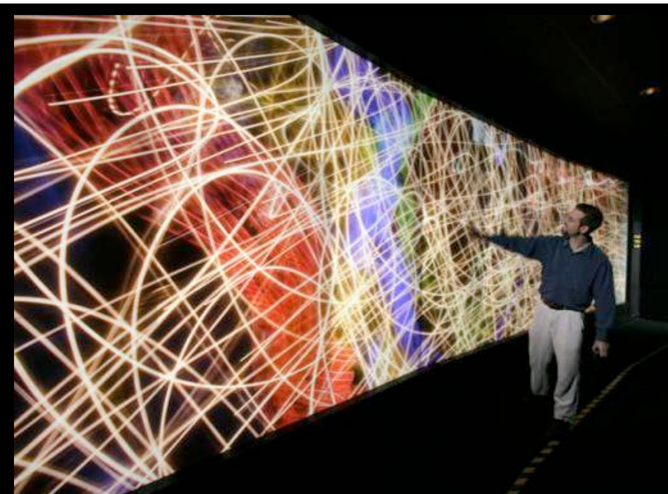

Scientific computing user support model

- "Whatever it takes" is the motto
- Share expertise in algorithms and application-development strategies
- Provide porting, tuning, optimization
- Help users in running applications, using application development tools and libraries
- Ensure application readiness by partnering with users to develop current next-generation applications
- Represent users' needs in OLCF planning and reporting exercises
  - Application requirements
  - Scientific progress and highlights
  - Issues with current resources

Expertise

The OLCF provides experts in user support, including PhD-level liaisons from fields such as chemistry, climate, physics, astrophysics, mathematics, numerical analysis, and computer science, who are also experts in developing code and optimizing it for the OLCF systems

Large projects are assigned liaisons to maximize opportunities for success on the leadership computing resources

Partnership with projects on LCF resources

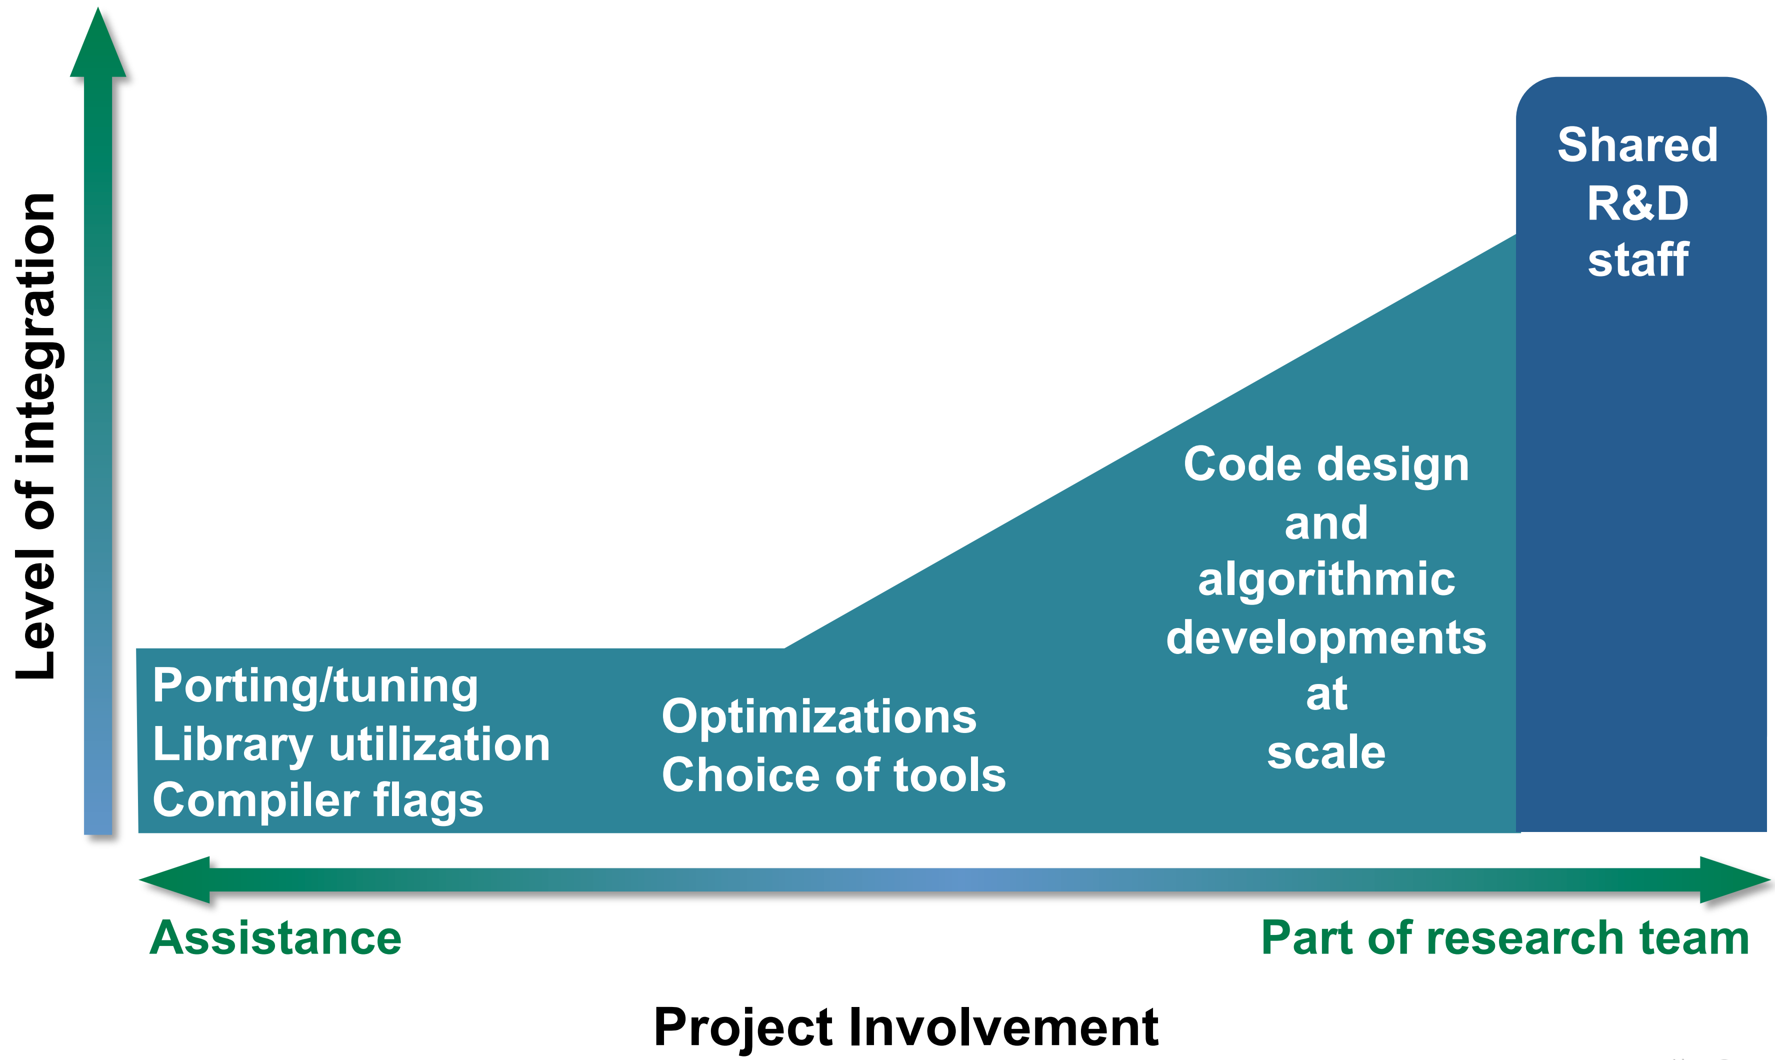

Liaison model helps maximize science

Project

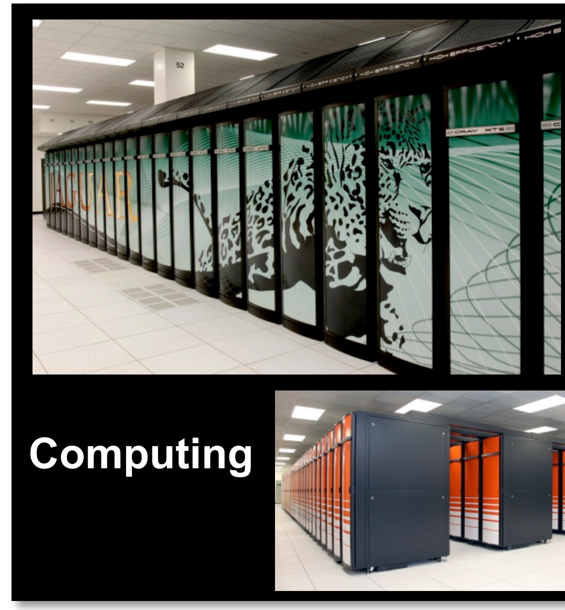

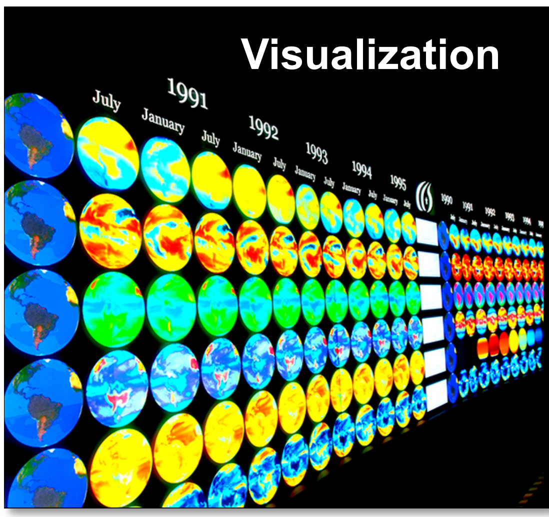

Whatever It Takes!

PI: Jihui Yang, General Motors

Materials and nanoscience

Nanostructural features in high-performance thermoelectric materials

Researchers simulate materials that can transform automobile waste heat directly into electricity

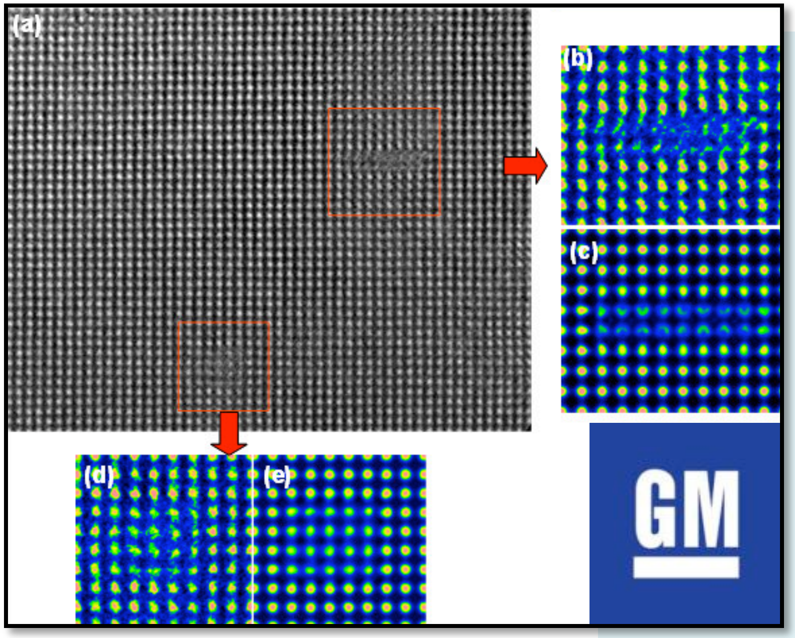

Of broad interest: The atomic structure of this material and the physical origin of its high thermoelectric nature are still under debate. This study will eventually explain these problems and lead to the discovery of new novel thermal-electric materials

- **Only 25% of auto fuel energy is used for vehicle mobility and accessories**
- **Team led by GM is working to develop waste heat recovery technology (a "thermoelectric alternator") for fuel-economy improvement**
- **Goal: 3–5% fuel-economy increase**
- **Largest-ever simulation—1728-atom supercell—made possible at LCF**
- **GM team identified the atomic configurations of the recently discovered $AgPb_mSbTe_{m+2}$ with DFT total energy calculations**
  - **Recent simulations of the $AgSbTe_2$ nanocluster embedded in the PbTe bulk material agree well with experimental measurements and open the way to explaining the origin of their outstanding thermal-electric performance**

OLCF liaison contributions

- **Worked with GM researchers to evolve a more scalable and capable version of VASP ab initio code**

PI: Warren Washington, NCAR

Climate
Modeling the full Earth system

New standard and low-emission climate scenarios at higher resolution

Final stage of ocean spin-up for full Earth System Model

100 years of ocean spin-up

20 years of carbon-cycle spin-up

Approaching steady state for sulfur cycle

Approaching release of CCSM4, the first Earth System Model

Simulated time evolution of the atmospheric CO_2 concentration originating from the land's surface

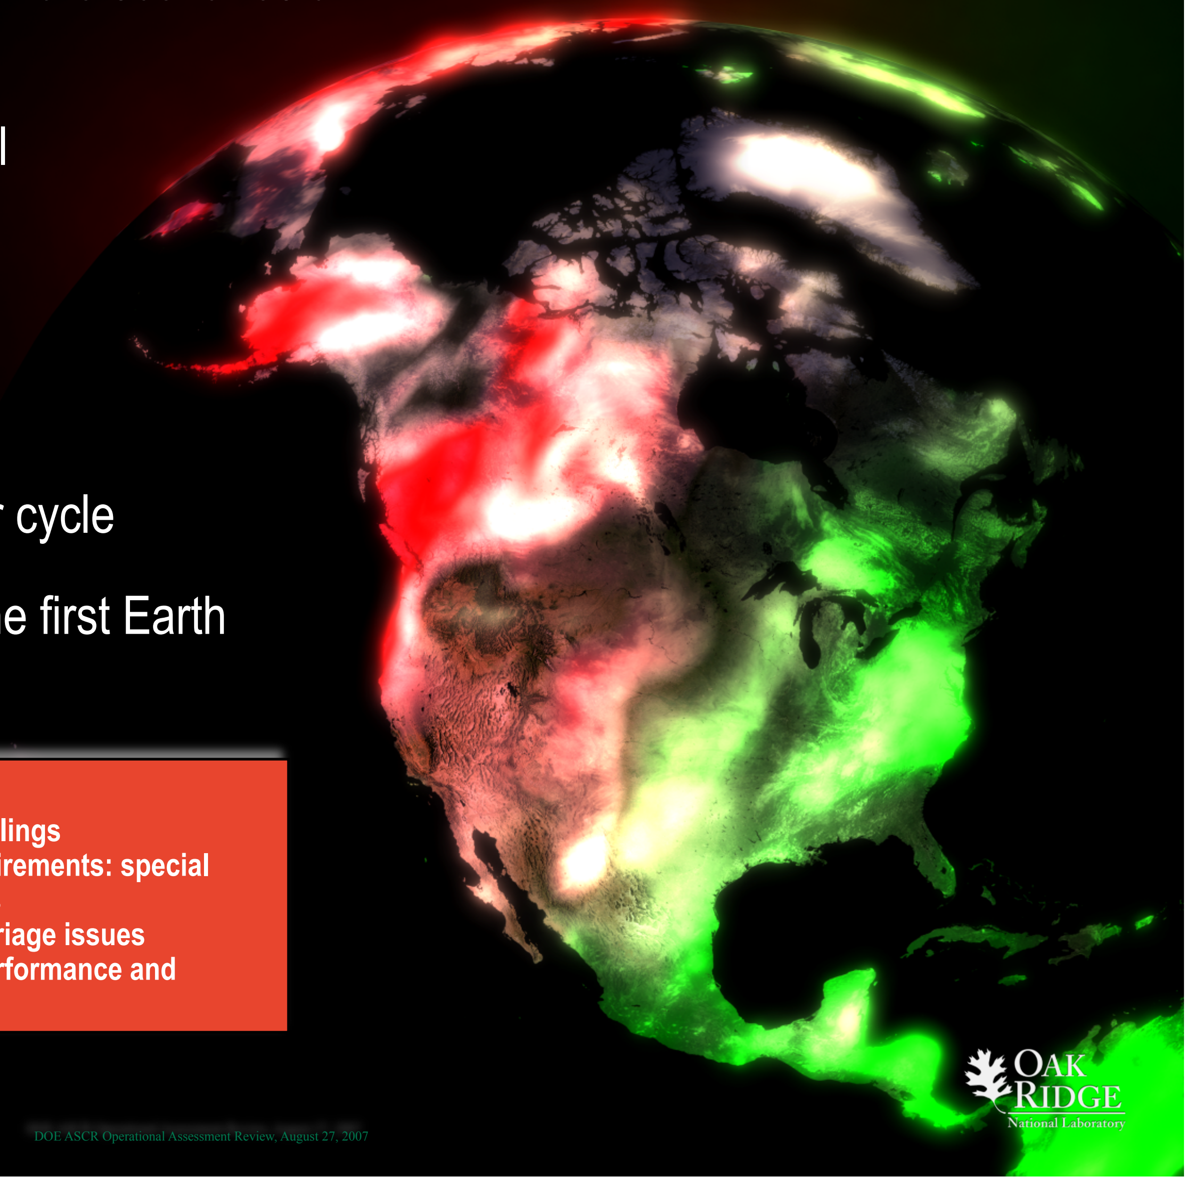

OLCF liaison contributions
- Advanced visualizations of new climate couplings
- Tuning of LCF infrastructure for climate requirements: special queues, reservations, dedicated file systems
- Weekly conference calls with developers to triage issues
- Development of timing libraries to assess performance and balance components

V. Balaji, NOAA/GFDL

Climate

Coupled high-resolution modeling of the Earth system

Science Objectives and Impact

- **Strategy:** The goal of decadal prediction requires long-time integrations of state-of-the-art models at unprecedented resolution
- **Driver:** What is the possibility of abrupt climate change on the timescale of decades rather than centuries?
- **Objective:** Use CM 2.5 for century simulations of the scientific issues associated with resolving mesoscale features in the atmospheric and ocean circulations, and its implications for understanding of forced/natural climate variability
- **Impact:** Provide reliable input for public policy relating to increasingly destructive storms in a warming world

Science Results

- CM2.5 vs CM2.1: increased ability to capture atmosphere-ocean covariability on mesoscale
- Recent CM2.5 runs on Jaguar/XT5 show agreement with observations of winds being stronger over warmer waters, weaker over cooler
- Finer-scale couple physics leading to qualitatively better depiction of tropical convection and climate in CM2.5
- **OLCF contribution:** Tuning and optimization, scheduling on resources (special queues)

Arabian sea winds and sea surface temperature
(Figure courtesy Whit Anderson, Rusty Benson, Tom Delworth, Tony Rosati and Gabe Vecchi, NOAA/GFDL)

The figure shows "the key feature of what increased resolution buys us: the ability to capture in a model atmosphere-ocean covariability on the mesoscale."
-V. Balaji, NOAA/GFDL

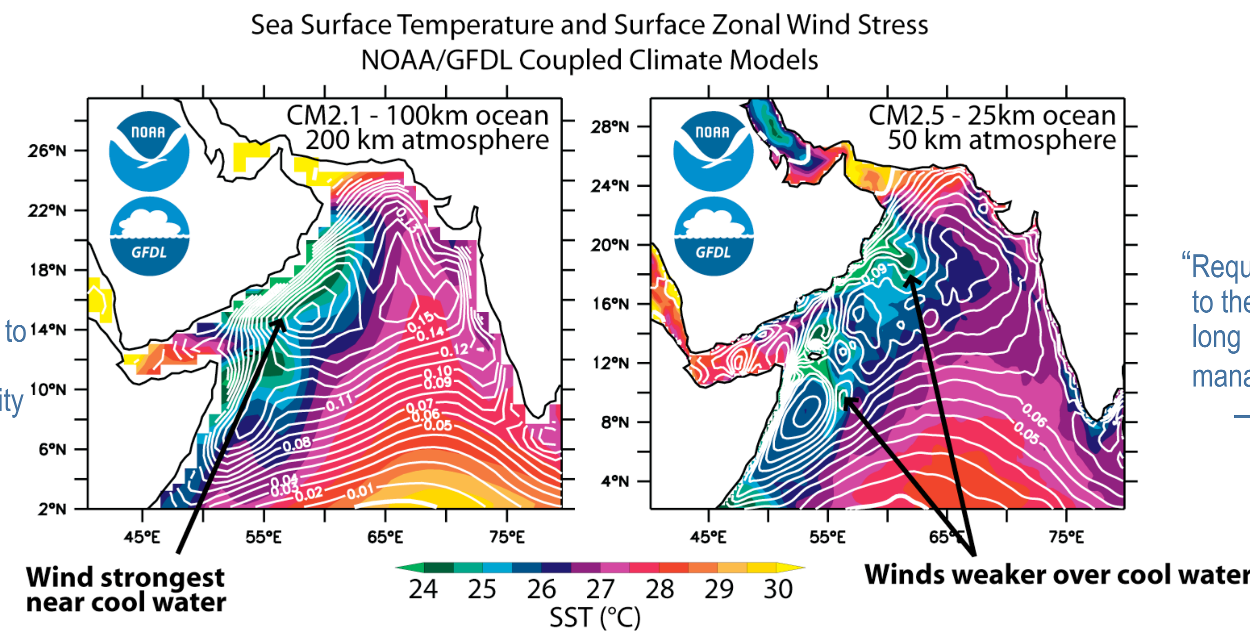

"Requires sustained access to the petascale for very long runs … currently being managed by a heroic effort"
–V. Balaji, NOAA/GFDL

PI: Fred Jaeger, ORNL

Fusion

Producing new insights for RF heating of ITER plasmas

Three-dimensional simulations of RF heating in the ITER fusion reactor as well as in present tokamaks (NSTX) shed new light on the behavior of superheated ionic gas

- When deuterium (D) or helium-3 (^{3}He) are used to damp the launched waves, they are accelerated to high energies, forming supra-thermal tails that significantly affect the wave propagation and absorption
- Energetic minority D ions enhance the fusion reaction rate
- Energetic ^{3}He ion tails form on both the tritium and ^{3}He distributions

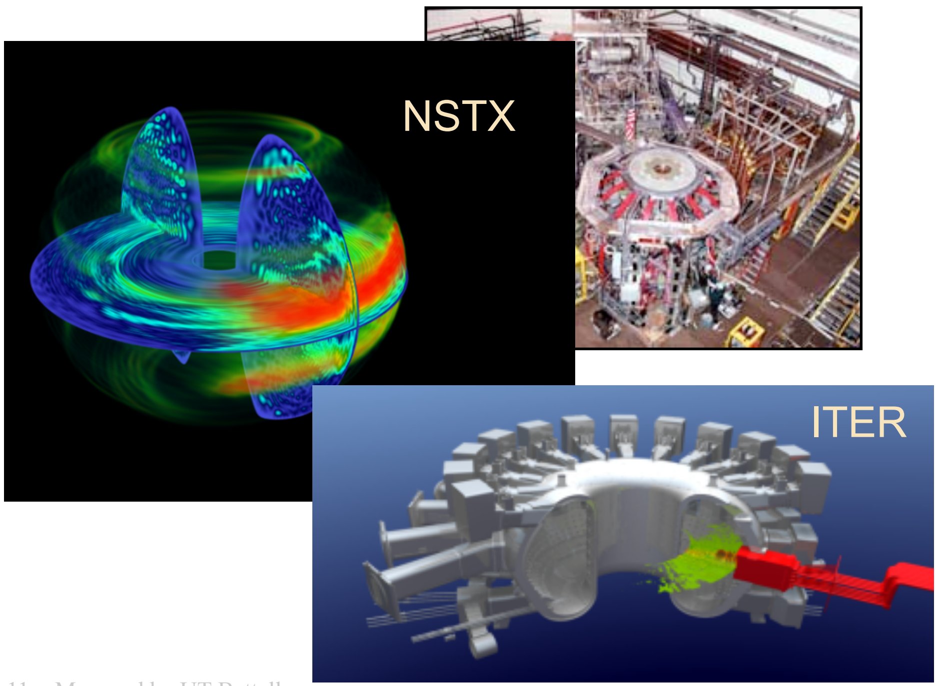

- **3-D simulations reveal new insights**
  - **"Hot spots" near antenna surface**
  - **"Parasitic" draining of heat to the plasma surface in smaller reactors**
- **Work pushing the boundaries of the system (28,900 cores, 154 TF) and demonstrating**
  - **Radial wave propagation and absorption**
  - **Efficient plasma heating**
- **AORSA's predictive capability can be coupled with Jaguar power to enhance fusion-reactor design and operation for an unlimited clean energy source**

OLCF liaison contributions

- Converted HPL from double real to double complex and replaced ScaLAPACK
- Acquired new version of BLAS from TACC
- Net performance gain of a factor of two

PI: Patrick Diamond, UCLA

Fusion

Breakthrough simulation of turbulent transport in ITER fusion plasmas

- Transition to favorable scaling of confinement for both ions and electrons now observed in GTC simulations for ITER plasmas
  - GTC simulation of electron turbulence used 28,000 cores for 42 hours producing 60 TB of data in a dedicated run on Jaguar as part of early-access 263 TF transition to operations
  - *[Yong (UCI), "Researchers Conduct Breakthrough Fusion Simulation,"* HPCWire *(July 14, 2008)]*
- GTC simulation critically tested quasilinear theory used in transport model
  - *[Z. Lin et al.,* Phys. Rev. Lett. *99, 265003 (December 2007)]*
- GTC simulation predicted properties of turbulent transport of fusion products
  - *[W. Zhang, Z. Lin, and L. Chen,* Phys. Rev. Lett. *101, in press]*
- GTC collaboration team
  - Led by Z. Lin of UC–Irvine and S. Klasky of ORNL

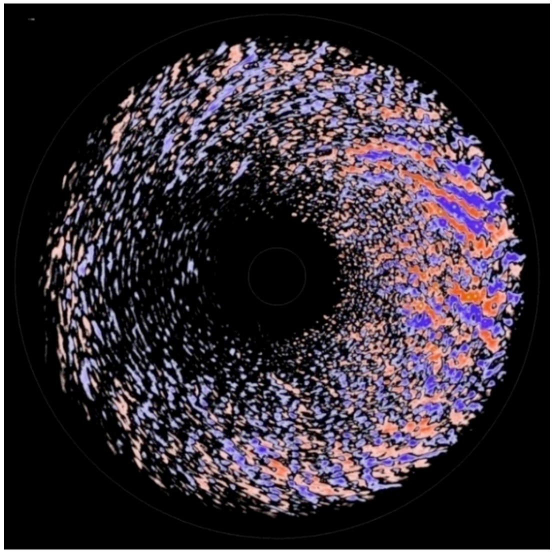

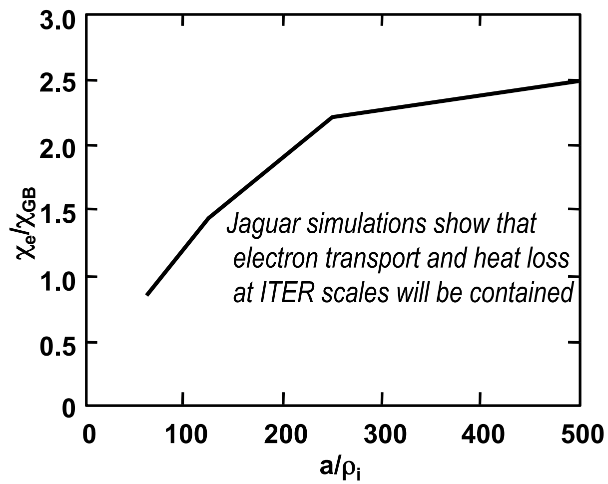

OLCF liaison contributions

- ADIOS for high-performance I/O. Wrote >60 TB of data at over 20 GB/second on 29K cores

Oak Ridge National Laboratory

Turbulent combustion

Jackie Chen, SNL

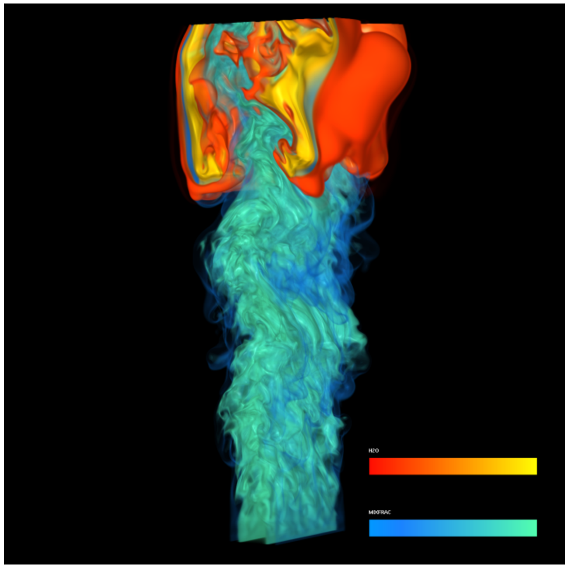

Science Objectives and Impact

- **Strategy:** Understand and model the effects of complex fuel chemistry in combustion engine conditions through high fidelity turbulence simulation
- **Driver:** Move to biofuel and coal-derived liquid fuels: Climate change and energy security
- **Objective:** Use of adaptive chemistry reduction and vectorization of new, larger chemistry models at low-temperature high pressure; efficient scaling to 140,000 cores
- **Impact:** Benchmark data for turbulent, igniting *n*-heptane and dimethyl ether flames at diesel engine pressures for model development and validation

Application Performance

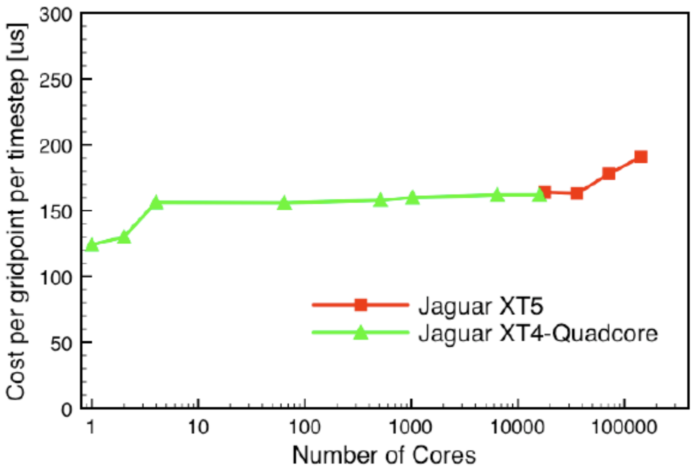

Science Results

- DNS of non-premixed *n*-heptane and dimethyl ether fuel jets issuing into high pressure, high temperature conditions to study the influence of multistage ignition chemistry on lifted flame stabilization under diesel thermochemical environment
- First 3-D simulation to fully resolve flame and ignition features including chemical composition, temperature profile, and turbulence flow characteristics

OLCF contribution

Application support and performance optimization for JaguarXT5 by R. Sankaran

Image caption

***n*-Heptane fuel jet (blue scale) and water vapor (red scale) indicating burned gas**

Geoscience

Modeling reactive flows in porous media

Peter Lichtner, LANL

Science Objectives and Impact

- **Strategy:** Evaluate remediation strategies for cleanup of uranium plume
- **Driver:** Potentially save billions of dollars in cleanup costs
- **Objective:** Perform high-fidelity simulations of the Hanford 300 Area and demonstrate scalability of PFLOTRAN to the petascale
- **Impact:** Stakeholders, DOE, and public are concerned that cleanup proceed to remediate contaminants along the Columbia River corridor —a resource for fish, recreation, and other uses

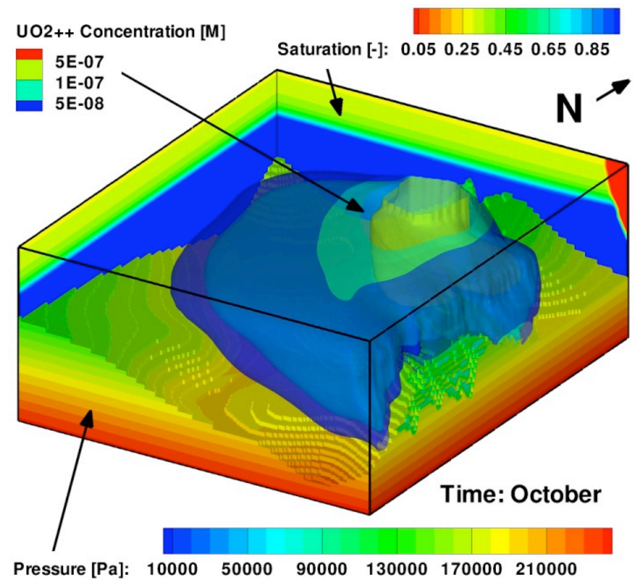

PFLOTRAN Performance

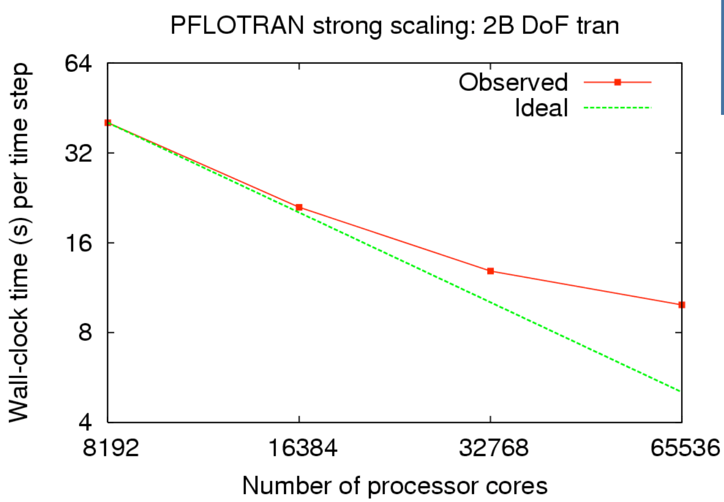

Science Results

- Demonstrated the ability of PFLOTRAN to scale to the petascale with complex nonlinear chemistry and highly transient fluid flow caused by fluctuations in the Columbia River stage
- Showed feasibility in using PFLOTRAN to model uranium dissipation at the Hanford 300 Area and evaluate remediation strategies

OLCF contribution

Ongoing development of multilevel solvers, parallel I/O performance optimization solver optimizations for full-machine runs on Jaguar

Image Caption (Top) **Isopleths of the simulated uranium plume at the Hanford 300 Area after 1 year. The Columbia River is on the right face. Also shown is the pressure and saturation fields.** (Left) **Departure from linear scaling is due to high cost of MPI_Allreduce() operations; performance should improve with a reformulated BiCGStab implementation with only one MPI_Allreduce() per iteration**

Preparing for the future

Application requirements: Process and actionable results

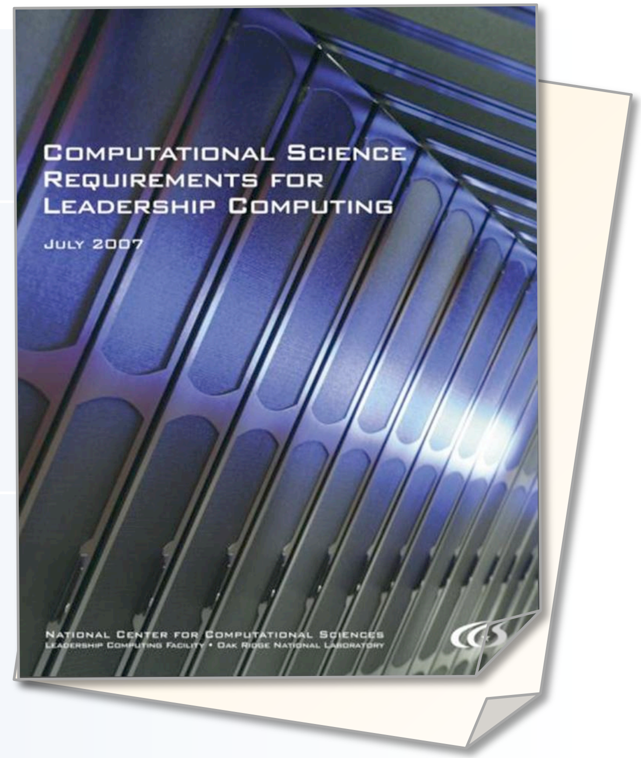

OLCF Application Requirements Council (ARC)	• Stood up in 2006 • Established ARC charter and requirements–management process
OLCF elicits requirements in many ways	• INCITE proposals • Questionnaires devised by LCF staff • One-on-one interviews • Existing publications/documentation • Analyzing source code
Application categories analyzed	• Science motivation and impact • Science quality and productivity • Application models, algorithms, software • Application footprint on platform • Data management and analysis • Early access science-at-scale scenarios
Results	• First annual >100-page application requirements document published September 2007 • New methods for categorizing platforms and application attributes devised and used in analysis: Guiding tactical infrastructure purchase and deployment • Best practice: Process being embraced and emulated by others

Innovation
Feedback loop for ensuring application readiness

Scientific Computing Group	Provides liaisons to application project teams
Application Requirements Council	Identifies application requirements
Technology Council	Decides how to best meet future application resource needs
Resource Utilization Council	Takes into account Science Team time constraints (e.g., upcoming meeting)

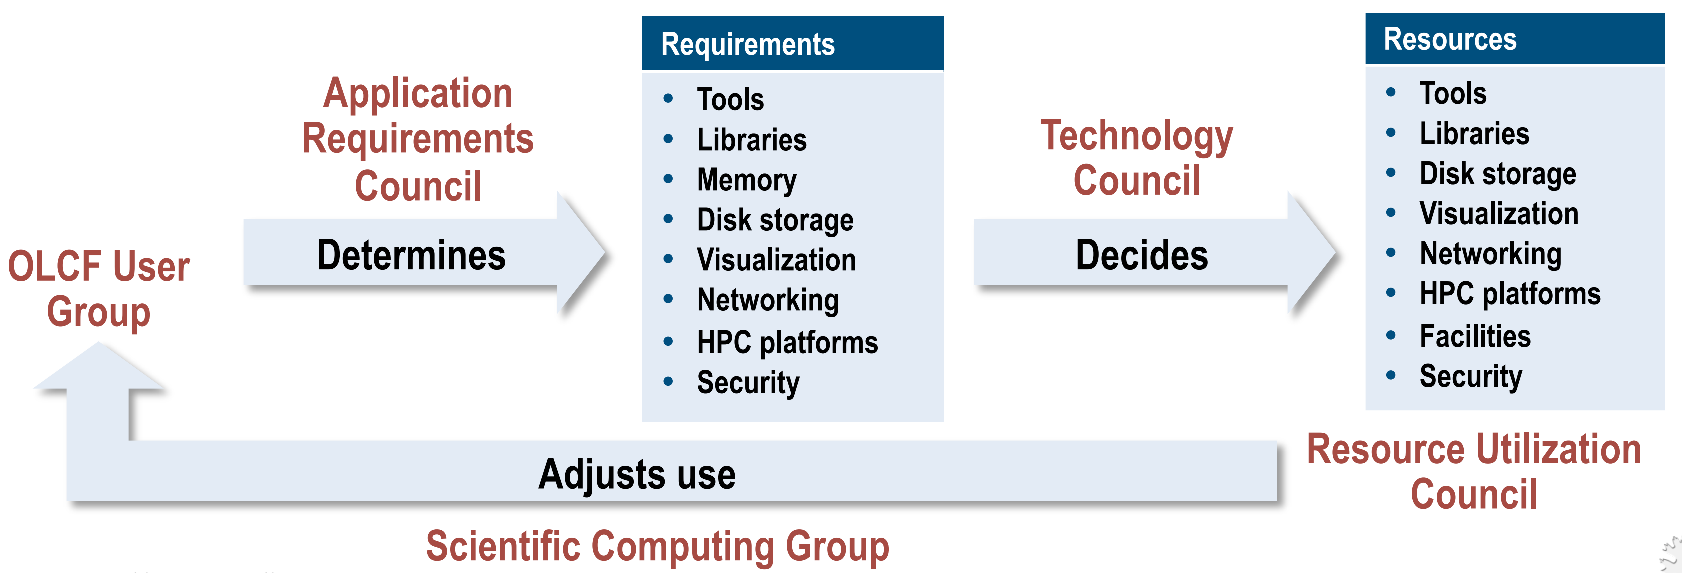

Software effectiveness on Jaguar/XT5
Improved computational science capabilities: Joule metric

- Four FY09 “Joule codes” were exercised extensively on Jaguar/XT5 during ES period
  - They possess science drivers for both strong and weak scaling
  - A very successful set of performance improvements have been realized for all four codes (by virtue of this metric)
- Raptor
  - Chemically reacting, turbulent multiphase flows in complex geometries
  - Combustion modeling of turbulent flame dynamics and mixing processes
  - Performance improvement: MPI halo communication redesign/refactor
- VisIt
  - Open source visualization and analysis tool for processing massive data sets
  - 2005 R&D 100 award winner with >100K downloads
  - Performance improvement: MPI all-to-all communication redesign/refactor
- CAM
  - Fifth generation Community Atmosphere Model of the NCAR AGCM used in climate studies
  - Performance improvement: moving to a hybrid MPI/OpenMP node model
- XGC1
  - 5-D gyrokinetic particle-in-cell code designed to model the whole plasma dynamics in experimentally realistic device geometry
  - Performance improvement: moving to a hybrid MPI/OpenMP node model
- Much more in the final FY09 Joule report
  - Take away: pushing scalability as driven by the science led to unanticipated performance improvements and fixes
  - Scaling is hard! Even the best people working on mature are confronted with surprises

Contact

Ricky A. Kendall

Scientific Computing
National Center for Computational Sciences
(865) 576-6905
kendallra@ornl.gov